

Operational Experience with the Frontier System in CMS

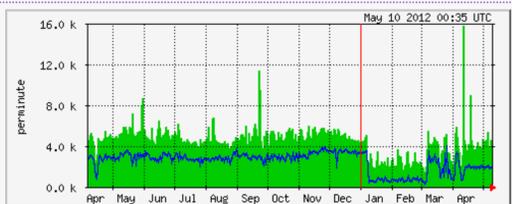


Barry Blumenfeld¹, Dave Dykstra², Peter Kreuzer³, Ran Du⁴, Weizhen Wang⁴

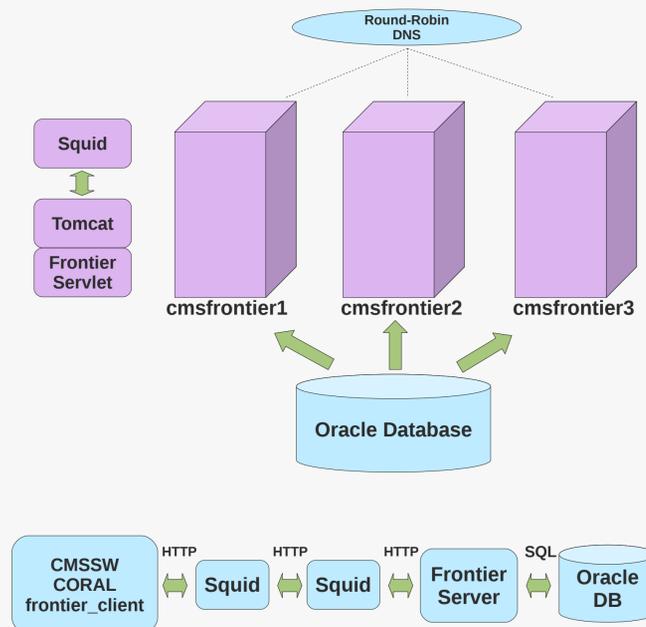
1: Johns Hopkins University (US), 2: Fermi National Accelerator Lab. (US)

3: Rheinisch-Westfaelische Tech. Hoch. (DE), 4: Chinese Academy of Sciences (CN)

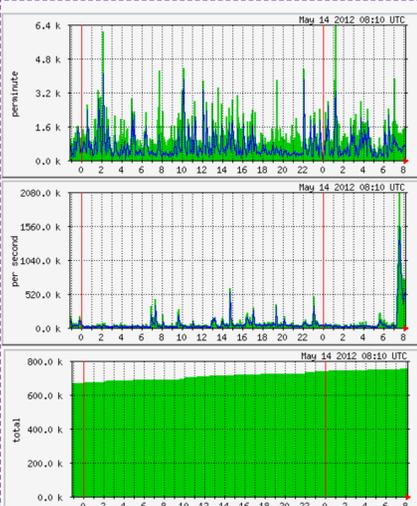
The **Frontier framework** is used in the CMS experiment at the LHC to deliver conditions data to processing clients worldwide, including calibration, alignment, and configuration information. Each of the central servers at CERN, called a **Frontier Launchpad**, uses **tomcat** as a **servlet** container to establish the communication between clients and the central **Oracle** database. HTTP-proxy **squid** servers, located close to clients, cache the responses to queries in order to provide high performance data access and to reduce the load on the central Oracle database. Each Frontier Launchpad also has its own reverse-proxy squid for caching.



The three central servers have been delivering about **5 million** responses every day since the LHC startup, containing about **40 GB** data in total, to more than **one hundred** Squid servers located worldwide, with an average response time on the order of **10 milliseconds**. The squid caches deployed worldwide process many more requests per day, over **700 million**, and deliver over **40 TB** of data. The chart above shows the load of Frontier Launchpad during the past year.



In order to guarantee the performance of the service and make the system easily maintainable, several monitoring tools to test the availability of remote squids, to analyze the tomcat log files and the accesses of the squid on the central Launchpad server have been developed.



MRTG: Multi Router Traffic Grapher

All the squids of CMS are monitored by an Open Source monitoring tool called MRTG which sends an SNMP request by UDP to every squid every 5 minutes. Because these charts are updated every 5 minutes they always contain the most up-to-date information available. There are three kinds of MRTG plots: Request/Fetch plots, In/Out plots, Object Count plots. Each kind of plots has four different lengths of duration displayed: Daily Graph, Weekly Graph, Monthly Graph and Yearly Graph.

In addition to the plots, each MRTG web page contains a useful message that shows the Squid version and the length of time it has been running.

Lpad_CH_CERN_2 Cache Statistics: HTTP Hits/Requests
System: cmsfrontier2
Maintainer: Barry Blumenfeld
Description: Frontier Server: Squid, cmsfrontier2
The statistics were last updated **Monday, 7 May 2012 at 23:45 UTC**.
At which time 'squid 2.7.STABLE9+fix28311+2833' had been up for 10 days, 9:34:32.

Hosts	Pages	Hits	Bytes	Last visit
vocms52.cern.ch	11547	11547	94.59 MB	16 May 2012 - 09:00
vocms58.cern.ch	100393	100393	84.62 MB	16 May 2012 - 08:58
vocms50.cern.ch	100388	100388	80.90 MB	16 May 2012 - 09:01
vocms61.cern.ch	71734	71734	89.46 MB	16 May 2012 - 09:00
cmscache01.grid.hep.ph.ic.ac.uk	19521	19521	58.56 MB	16 May 2012 - 09:00
higgs.fr.cinvestav.mx	8571	8571	11.18 MB	16 May 2012 - 08:58
cmsrv35.hal.gov	7928	7928	33.18 MB	16 May 2012 - 09:00
cmsrv79.hal.gov	7625	7625	33.96 MB	16 May 2012 - 08:53
cmsrv78.hal.gov	6808	6808	29.82 MB	16 May 2012 - 09:00
capicm.hep.wisc.edu	5902	5902	8.44 MB	16 May 2012 - 09:01
cmsrv34.hal.gov	5625	5625	33.25 MB	16 May 2012 - 09:01
chep250.ultraflight.org	5415	5415	9.05 MB	16 May 2012 - 08:52
red-squid2.uni.edu	4973	4973	5.49 MB	16 May 2012 - 08:52
ingrid-monbox.cism.ucl.ac.be	4524	4524	2.54 MB	16 May 2012 - 09:00
cayenne.hep.wisc.edu	4475	4475	2.08 MB	16 May 2012 - 08:52
hs1.accre.vanderbilt.edu	4356	4356	65.47 MB	16 May 2012 - 08:57
cmsosquid.psu.fr	4268	4268	5.36 MB	16 May 2012 - 08:53
red-squid1.uni.edu	4225	4225	2.87 MB	16 May 2012 - 08:50
proxy-1.2.cd.edu	4133	4133	2.56 MB	16 May 2012 - 08:56

AWStats

AWStats is an Open Source program which shows detailed information from the log files of the squids in each Frontier Launchpad. This program updates every hour. In addition to being useful by itself, it is the basis for the Non-Proxy Monitoring Tools.

Squid not Responding List	
T3_OR_IASA	Friday, 11 May 2012 at 14:15 UTC
T3_PL_Warsaw	Friday, 11 May 2012 at 14:15 UTC
T3_US_Catheth	Friday, 11 May 2012 at 14:15 UTC
T3_US_FIT	Friday, 11 May 2012 at 14:15 UTC
T3_US_NotreDome	Friday, 11 May 2012 at 14:15 UTC

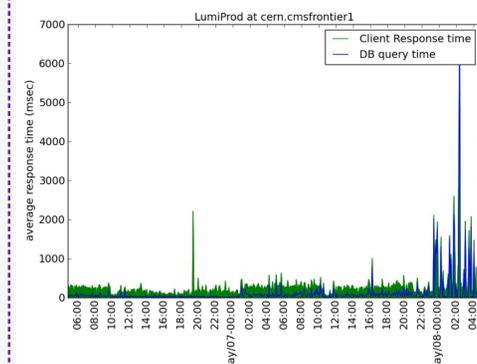
Squid Not Responding List

The Squid Not Responding List monitoring tool is used to find out which squids are not currently active. The activeness of one squid is tested by filtering the corresponding MRTG plots web page; all the squids whose MRTG web page doesn't contain 'at which time' will be seen as not-responding squids and added to the not-responding list.

SAM Test

There are Squid and Frontier tests in the Site Availability Monitor (SAM). These tests run every hour and are used to test the availability of each site, but sometimes have a large latency.

Status	NA	OK	WARNING	CRITICAL	UNKNOWN	MISSING	MAINTENANCE	REMOVED
Legend:	NA	OK	W	C	U	M	MT	RM

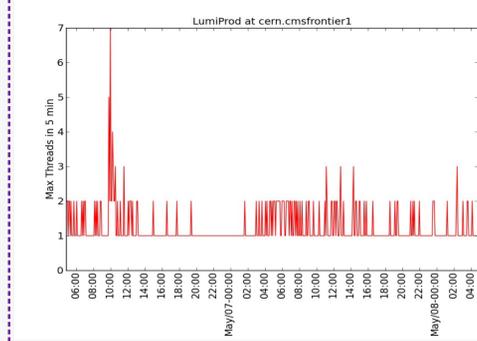


Max Threads Monitoring Tool

The number of threads refers to the number of concurrent requests to one servlet on a Frontier Launchpad. Each servlet has its own limitation of the maximum amount of concurrent requests.

The Max Thread Monitoring Tool was developed to monitor the number of requests to each servlet on Frontier Launchpad. Whenever the number of threads goes over 75% of the maximum allowed for any servlet, an alarm message is emailed to the operations team so they can investigate the cause.

At the same time, because DB response time has a significant effect on the number of threads, the Max Threads Monitoring Tool also monitors DB response time.

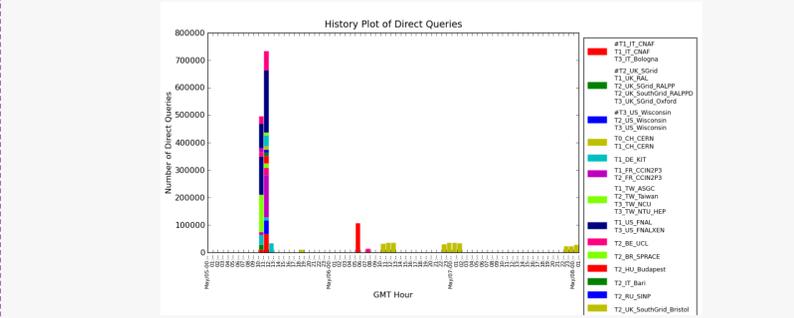


Squid Src Compare

Site	Match/NOT	geolist.txt	mrtg
TO CH_CERN	OK	+cmsst0frontier1.cern.ch; +cmsst0frontier2.cern.ch; +cmsst0frontier3.cern.ch; +cmsst0frontier4.cern.ch;	cmsst0frontier1.cern.ch; cmsst0frontier2.cern.ch; cmsst0frontier3.cern.ch; cmsst0frontier4.cern.ch;
T1 CH_CERN	OK	cmsst0frontier.cern.ch; cmsst0frontier1.cern.ch; cmsst0frontier2.cern.ch;	cmsst0frontier.cern.ch; cmsst0frontier1.cern.ch; cmsst0frontier2.cern.ch;
T1 DE_KIT	OK	cmsstq1-fzk.gridka.de; cmsstq2-fzk.gridka.de;	cmsstq1-fzk.gridka.de; cmsstq2-fzk.gridka.de;
T1 ES_PIC	OK	squid01.pic.es; squid02.pic.es; squid03.pic.es; squid04.pic.es;	squid01.pic.es; squid02.pic.es; squid03.pic.es; squid04.pic.es;
T1 FR_CERN2P3	OK	cdlegms01.m2p3.fr; cdlegms02.m2p3.fr;	cdlegms01.m2p3.fr; cdlegms02.m2p3.fr;
T1 IT_CNAF	OK	squid-cms-1.ccr.cnaf.infn.it; squid-cms-3.ccr.cnaf.infn.it;	squid-cms-1.ccr.cnaf.infn.it; squid-cms-3.ccr.cnaf.infn.it;
T1_TW_ASCG	OK	lsg00114.grid.sinica.edu.tw; squid01.grid.sinica.edu.tw; squid02.grid.sinica.edu.tw;	lsg00114.grid.sinica.edu.tw; squid01.grid.sinica.edu.tw; squid02.grid.sinica.edu.tw;
T1 UK_RAL	OK	lsgsquid01.gridpp.rl.ac.uk; lsgsquid02.gridpp.rl.ac.uk;	lsgsquid01.gridpp.rl.ac.uk; lsgsquid02.gridpp.rl.ac.uk;
T1 US_FNAL	OK	cmsst0frontier1.fnal.gov; cmsst0frontier2.fnal.gov; cmsst0frontier3.fnal.gov; cmsst0frontier4.fnal.gov;	cmsst0frontier1.fnal.gov; cmsst0frontier2.fnal.gov; cmsst0frontier3.fnal.gov; cmsst0frontier4.fnal.gov;
T2 AT_Vienna	OK	cmsgrid.oew.ac.at;	cmsgrid.oew.ac.at;
T2 BE_IIIHE	OK	frontier.ihc.ac.be;	frontier.ihc.ac.be;
T2 BE_UCL	OK	ingrid-monbox.cism.ucl.ac.be;	ingrid-monbox.cism.ucl.ac.be;
T2 BR_SPRACE	OK	frontier.sprace.org.br;	frontier.sprace.org.br;
T2 BR_URJ	OK	+phedex.heppgrid.uerj.br;	phedex.heppgrid.uerj.br;
T2 CH_CAF	OK	cmsst0frontier.cern.ch; cmsst0frontier1.cern.ch; cmsst0frontier2.cern.ch;	cmsst0frontier.cern.ch; cmsst0frontier1.cern.ch; cmsst0frontier2.cern.ch;
T2 CH_CSCS	OK	cmsvobox.lcg.cscs.ch;	cmsvobox.lcg.cscs.ch;
T2 CN_Beijing	OK	vobox.lhep.ac.cn;	vobox.lhep.ac.cn;
T2 DE_DESY	OK	T2-cms-dbl.desy.de; T2-cms-dbl.desy.de;	T2-cms-dbl.desy.de; T2-cms-dbl.desy.de;
T2 DE_RWTH	OK	grid-squid1.physik.rwth-aachen.de; grid-squid2.physik.rwth-aachen.de;	grid-squid1.physik.rwth-aachen.de; grid-squid2.physik.rwth-aachen.de;
T2 EE_Estonia	OK	io.hep.kit.fi.ee;	io.hep.kit.fi.ee;
T2 ES_CIEMAT	OK	lscgca01.ciemat.es;	lscgca01.ciemat.es;
T2 ES_IPCA	OK	squid01.prv.ica.es; squid02.prv.ica.es; +squid01.ica.es; +squid02.ica.es;	squid01.prv.ica.es; squid02.prv.ica.es; +squid01.ica.es; +squid02.ica.es;

Squid Source Compare Monitoring Tool

Sites that get conditions data from Frontier are supposed to register their Squid(s) in CVS. At the same time, all Squids are monitored by MRTG. Sometimes the Squid information in CVS becomes inconsistent with the information in MRTG (e.g. an administrator doesn't keep CVS up to date, Squids are shared by multiple sites, or Squids have different addresses on private and public networks). In order to make sure the information stays as correct as possible, the Squid Source Comparing Monitoring tool was developed to keep Squid information consistent between CVS and MRTG. It sends an email warning whenever the two sources become different, and it includes the ability to store exceptions for cases where the operations team understands why the two sources are different.



Non-Proxy Monitoring Tools

The Non-proxy problem refers to the possibility that jobs bypass the local Squid and directly access the Frontier Launchpads, which can cause a heavy load on the Frontier Launchpad system and low efficiency. In order to detect such problems, two monitoring tools were developed, one is called "Failover nodes", the other is called "Failover summary". The above chart shows the latter one. The former tool generates a web log and automatically sends an email warning to the administrator of a Squid when jobs at the Squid's site directly access the Frontier Launchpads too many times in an hour.